On the Computation of Distances for Probabilistic Context-Free Grammars

Colin de la Higuera James Scicluna Mark-Jan Nederhof

LINA, UMR 6241
University of Nantes*
France
cdlh@univ-nantes.fr james.scicluna@univ-nantes.fr

School of Computer Science
University of St Andrews
UK

arXiv:1407.1513v1 [cs.FL] 6 Jul 2014

Probabilistic context-free grammars (PCFGs) are used to define distributions over strings, and are powerful modelling tools in a number of areas, including natural language processing, software engineering, model checking, bio-informatics, and pattern recognition. A common important question is that of comparing the distributions generated or modelled by these grammars: this is done through checking language equivalence and computing distances. Two PCFGs are language equivalent if every string has identical probability with both grammars. This also means that the distance (whichever norm is used) is null. It is known that the language equivalence problem is interreducible with that of multiple ambiguity for context-free grammars, a long-standing open question. In this work, we prove that computing distances corresponds to solving undecidable questions: this is the case for the L1, L2 norm, the variation distance and the Kullback-Leibler divergence. Two more results are less negative: 1. The most probable string can be computed, and, 2. The Chebyshev distance (where the distance between two distributions is the maximum difference of probabilities over all strings) is interreducible with the language equivalence problem.

1 General motivation and introduction

Probabilistic context-free grammars (PCFGs) are powerful modelling tools in a number of areas, including natural language processing, software engineering, model checking, bio-informatics, and pattern recognition. In natural language processing, these grammars are used as language models [23, 3] or for parsing natural language [22, 24]. In model checking the crucial questions of program equivalence or meeting specifications will often be solved through tackling the grammar equivalence problem [9, 11].

In pattern recognition, probabilistic context-free grammars have been proposed and used for 40 years [13]. In bio-informatics structural dependencies are modelled through context-free rules, whose probabilities can be estimated [32, 33].

In all these areas, the following questions are important: given two grammars, are they equivalent? Two grammars are equivalent (*strongly*, or *language equivalent*) when every string has identical probability in each distribution. More generally, a distance between distributions expresses how close they are, with a zero distance coinciding with equivalence.

Furthermore, in many areas, these probabilistic models are to be learnt. When learning, comparison between states or nonterminals often determines if a merge or a generalization is to take place. Key grammatical inference operations [16] will depend on the precision with which an equivalence is tested or a distance is measured.

In the case of probabilistic finite automata, these questions have been analysed with care. The initial study by Balasubramanian [2] showed that the equivalence problem for hidden Markov models admitted

*The first and second author acknowledge partial support by the Région des Pays de la Loire.

Submitted to:

a polynomial algorithm. Later, a number of authors showed that the Euclidian distance could be computed in polynomial time [28]. This important result made use of the key concept of co-emission. A similar result was obtained for PFAs [29].

Negative results were also obtained: The L1 distance, and the variation distance were proved to be intractable [26]. Relative entropy (or Kullback-Leibler divergence) cannot be computed between general PFAs. But it can be computed for deterministic [4] and unambiguous [8] PFAs.

A related problem, that of computing the most probable string (also called the *consensus string*) was proved to be an NP-hard problem [5, 27]; heuristics and parameterized algorithms have been proposed [17]. Some of these results and techniques were also extended to PCFGs [18].

The co-emission of a PCFG and a hidden Markov model was discussed by [19], who formulated the problem in terms of finding the solution to a system of quadratic equations; for the difficulties in solving such systems, see also [10]. By a related construction, a probabilistic finite automaton can be obtained from a given unambiguous (non-probabilistic) finite automaton and a given PCFG, in such a way that the Kullback-Leibler divergence between the two probability distributions is minimal [30].

The same problems have been far less studied for PCFGs: the equivalence problem has recently been proved [12] to be as hard as the multiple ambiguity problem for context-free grammars: do two context-free grammars have the same number of derivation trees for each string? Since on the other hand it is known [1] that probabilistic pushdown automata are equivalent to PCFGs, it follows that the equivalence problem is also interreducible with the multiple ambiguity problem. Before that, the difficulty of co-emission and of related results was shown in [19].

Since computing distances is at least as hard as checking equivalence (the case where the distance is 0 giving us the answer to the equivalence question) it remains to be seen, for PCFGs, just how hard it is to compute a distance, and also to see if all distances are as hard.

This is the subject of this work. We have studied a number of distances over distributions, including the L1 and L2 norms, the Hellinger and the variation distances and the Kullback-Leibler divergence. We report that none of these can be computed.

On the other hand, the fact that the consensus string can be computed allows to show that the Chebyshev distance (or L_∞ norm) belongs to the same class as the multiple ambiguity problem for context-free grammars and the equivalence problem for PCFGs.

In Section 2 we remind the reader of the different notations, definitions and key results we will be needing. In Section 3 we go through the new results we have proved. We conclude in Section 4.

2 Definitions and notations

Let $[n]$ denote the set $\{1,\ldots,n\}$ for each $n \in \mathbb{N}$. Logarithms will be taken in base 2. An *alphabet* Σ is a finite non-empty set of symbols. A *string* w over Σ is a finite sequence $w = a_1 \ldots a_n$ of symbols. Symbols will be indicated by $a, b, c, \ldots$, and strings by $u, v, \ldots, z$. Let $|w|$ denote the length of w. The *empty string* is denoted by λ.

We denote by $\Sigma^\star$ the set of all strings and by $\Sigma^{\leq n}$ the set of those of length at most n. Similarly, $\Sigma^n = \{x \in \Sigma^\star : |x| = n\}$, $\Sigma^{<n} = \{x \in \Sigma^\star : |x| < n\}$ and $\Sigma^{\geq n} = \{x \in \Sigma^\star : |x| \geq n\}$.

A *probabilistic language* $\mathscr{D}$ is a probability distribution over $\Sigma^\star$. The probability of a string $x \in \Sigma^\star$ under the distribution $\mathscr{D}$ is denoted as $Pr_{\mathscr{D}}(x)$ and must satisfy $\sum_{x\in\Sigma^\star} Pr_{\mathscr{D}}(x) = 1$. If A is a language (thus a set of strings, included in $\Sigma^\star$), and $\mathscr{D}$ a distribution over $\Sigma^\star$, $Pr_{\mathscr{D}}(A) = \sum_{x\in A} Pr_{\mathscr{D}}(x)$.

If the distribution is modelled by some grammar G, the probability of x according to the probability distribution defined by G is denoted by $Pr_G(x)$. The distribution modelled by a grammar G will be

denoted by $\mathscr{D}_G$.

2.1 Context-free grammars

A *context-free grammar* is a tuple $<\Sigma, V, R, S>$ where Σ is a finite alphabet (of terminal symbols), V is a finite alphabet (of variables or non-terminals), $R \subset V \times (\Sigma \cup V)^*$ is a finite set of production rules, and S $(\in V)$ is the axiom or start symbol.

We will write $N \to \beta$ for rule $(N, \beta) \in R$. If $\alpha, \beta, \gamma \in (\Sigma \cup V)^*$ and $(N, \beta) \in R$ we have $\alpha N \gamma \Rightarrow \alpha \beta \gamma$. This means that string $\alpha N \gamma$ *derives* (in one step) into string $\alpha \beta \gamma$.

$\stackrel{*}{\Longrightarrow}$ is the reflexive and transitive closure of $\Rightarrow$. If there exists $\alpha_0, \ldots, \alpha_k$ such that $\alpha_0 \Rightarrow \cdots \Rightarrow \alpha_k$ we will write $\alpha_0 \stackrel{k}{\Rightarrow} \alpha_k$. $\mathbb{L}(G)$ is the language generated by G: the set of all strings w over Σ such that $S \stackrel{*}{\Longrightarrow} w$.

A sequence $\alpha_0 \Rightarrow \cdots \Rightarrow \alpha_k$ is a derivation of α_k from α_0. A derivation step $\alpha N \gamma \Rightarrow \alpha \beta \gamma$ is a *left-most derivation step* if $\alpha \in \Sigma^*$. A derivation is *left-most* if each step is left-most.

A context-free grammar is *proper* if it satisfies the following three properties:

1. It is cycle-free, i.e. no non-terminal A exists such that $A \stackrel{+}{\Longrightarrow} A$.
2. It is λ-free, i.e. either no rules with λ on the RHS exist or exactly one exists with S on the LHS (i.e. $S \to \lambda$) and S does not appear on the RHS of any other rule.
3. It contains no useless symbols or non-terminals. This means that every symbol and non-terminal should be reachable from S and every non-terminal should derive at least one string from Σ^*.

A context-free grammar is ambiguous if there exists a string w admitting two different left-most derivations from S to w. Given any string w, we can define the multiplicity $m_G(w)$ as the number of different left-most derivations from S to w. If $\forall w \in \Sigma^\star\ m_G(w) \leq 1$, G is unambiguous. Otherwise it is ambiguous. If $\forall w \in \Sigma^\star\ m_G(w) < \infty$, G is a *finite multiplicity grammar*. If a grammar is proper it has finite multiplicity. Two finite multiplicity grammars G_1 and G_2 are multiplicity equivalent if $\forall w \in \Sigma^\star\ m_{G_1}(w) = m_{G_2}(w)$.

The multiplicity equivalence problem has been studied for many years [25]: the problem has been proved to be decidable only for particular classes of grammars.

Results regarding the decidability of context-free grammars can be found in many textbooks [15]:

1. Given two context-free grammars G_1 and G_2, the (*equivalence*) question $\mathbb{L}(G_1) = \mathbb{L}(G_2)$? is undecidable.
2. Given two context-free grammars G_1 and G_2, the (*inclusion*) question $\mathbb{L}(G_1) \subseteq \mathbb{L}(G_2)$? is undecidable.
3. Given two context-free grammars G_1 and G_2, the (*emptiness of intersection*) question $\mathbb{L}(G_1) \cap \mathbb{L}(G_2) = \emptyset$? is undecidable.

2.2 Probabilistic context-free grammars

Definition 1. *A* probabilistic context-free grammar (PCFG) *G is a context-free grammar $<\Sigma, V, R, S>$ with a probability function $P : R \to \mathbb{R}^+$.*

$Pr_G(x)$ is the sum of all the leftmost derivations' probabilities of x, where the probability of a leftmost derivation is the product of the rule probabilities used in the derivation. A PCFG G is said to be consistent

if $\sum_{x\in\Sigma^*} Pr_G(x) = 1$. Unless otherwise specified, any PCFG mentioned from now onwards is assumed to be consistent.

Parsing with a PCFG is usually done by adapting the Earley or the CKY algorithms [21]. By straightforward variants allowing every terminal to match every input position, one can compute $Pr_G(\Sigma^n)$, still in polynomial time in n. By summing $Pr_G(\Sigma^i)$ for $i < n$ one obtains $Pr_G(\Sigma^{<n})$, and $Pr_G(\Sigma^{\geq n})$ is $1 - Pr_G(\Sigma^{<n})$. Alternatively, $Pr_G(\Sigma^{\geq n})$ can be computed directly by variants of algorithms computing prefix probabilities [20, 34].

We denote by $\mathbb{L}(G)$ the support language of G, ie the set of strings of non null probability. The class of all PCFGs over alphabet Σ will be denoted by $\mathbf{PCFG}(\Sigma)$.

There exists an effective procedure which, given a proper CFG G, builds a PCFG G' such that $\forall x \in \Sigma^*, Pr_{G'}(x) > 0 \iff x \in \mathbb{L}(G)$. We call this procedure MP for Make Probabilistic.

One possible procedure for MP is to first assign uniform probabilities to the given CFG, thus obtaining a possibly inconsistent PCFG which then can be converted into a consistent PCFG using the procedure explained in [14].

Let us formally define the (language) equivalence problem:

Definition 2. *Two probabilistic grammars G_1 and G_2 are* (language) equivalent *if $\forall x \in \Sigma^*, Pr_{G_1}(x) = Pr_{G_2}(x)$. We denote by $\langle EQ, \mathbf{PCFG}(\Sigma)\rangle$ the decision problem: are two PCFGs G_1 and G_2 equivalent?*

The following result holds for probabilistic pushdown automata, which are shown in [1] to be equivalent to PCFGs.

Proposition 1. *[12] The $\langle EQ, \mathbf{PCFG}(\Sigma)\rangle$ problem is interreducible with the multiplicity equivalence problem for CFGs.*

2.3 About co-emissions

Co-emission has been identified as a key concept allowing, in the case of hidden Markov models or probabilistic finite-state automata, computation in polynomial time of the distance for the L2 norm (and more generally any Lp norm, for even values of p): the distance can be computed as a finite sum of co-emissions.

Definition 3. *The coemission of two probabilistic grammars G_1 and G_2 is the probability that both grammars G_1 and G_2 simultaneously emit the same string:*
$\text{COEM}(G_1, G_2) = \sum_{x\in\Sigma^*} Pr_{G_1}(x) \cdot Pr_{G_2}(x)$

A particular case of interest is the probability of twice generating the same string when using the same grammar: Given a PCFG G, the *auto-coemission* of G, denoted as $\text{AC}(G)$, is $\text{COEM}(G, G)$. If the grammars are ambiguous, internal factorization is required in the computation of co-emission. In order to detect this we introduce the *tree-auto-coemission* as the probability that the grammar produces exactly the same left-most derivation (which corresponds to a specific tree):

Definition 4. *Given a* PCFG *G, the* tree-auto-coemission *of G is the probability that G generates the exact same left-most derivation twice. We denote it by* $\text{TAC}(G)$.

Note that the tree-auto-coemission and the auto-coemission coincide if and only if G is unambiguous:

Proposition 2. *Let G be a* PCFG.
$\text{AC}(G) \geq \text{TAC}(G)$.
$\text{AC}(G) = \text{TAC}(G) \iff$ *G is unambiguous.*

2.4 Distances between distributions

A PCFG defines a distribution over $\Sigma^\star$. If two grammars can be compared syntactically, they can also be compared semantically: do they define identical distributions, and, if not, how different are these distributions?

Definition 5. *The* L1 *distance (or Manhattan distance) between two probabilistic grammars G_1 and G_2 is:*

$$d_{\mathrm{L1}}(G_1,G_2)=\sum_{x\in\Sigma^*}|Pr_{G_1}(x)-Pr_{G_2}(x)|$$

Definition 6. *The* L2 *distance (or Euclidian distance) between two probabilistic grammars G_1 and G_2 is:*

$$d_{\mathrm{L2}}(G_1,G_2)=\sqrt{\sum_{x\in\Sigma^*}(Pr_{G_1}(x)-Pr_{G_2}(x))^2}$$

L2 distance can be rewritten in terms of coemission, as:

$$d_{\mathrm{L2}}(G_1,G_2)=\sqrt{\textsc{CoEm}(G_1,G_1)-2\textsc{CoEm}(G_1,G_2)+\textsc{CoEm}(G_2,G_2)}$$

Definition 7. *The* L_∞ *distance (or Chebyshev distance) between two probabilistic grammars G_1 and G_2 is:*

$$d_{\mathrm{L}_\infty}(G_1,G_2)=\max_{x\in\Sigma^*}|Pr_{G_1}(x)-Pr_{G_2}(x)|$$

Note that the L_∞ distance seems closely linked with the consensus string, which is the most probable string in a language.

Definition 8. *The* variation distance *between two probabilistic grammars G_1 and G_2 is:*

$$d_V(G_1,G_2)=\max_{X\subset\Sigma^*}\sum_{x\in X}(Pr_{G_1}(x)-Pr_{G_2}(x))$$

The variation distance looks like d_{L_∞}, but is actually connected with d_{L1}:

$$d_V(G_1,G_2)=\frac{1}{2}d_{\mathrm{L1}}(G_1,G_2) \tag{1}$$

A number of distances have been studied elsewhere (for example [19, 7]):

Definition 9. *The* Hellinger distance *between two probabilistic grammars G_1 and G_2 is:*

$$d_H(G_1,G_2)=\frac{1}{2}\cdot\sum_{w\in\Sigma^*}\left(\sqrt{Pr}_{G_1}(w)-\sqrt{Pr}_{G_2}(w)\right)^2$$

The Jensen-Shannon (JS) *distance between two probabilistic grammars G_1 and G_2 is:*

$$d_{JS}(G_1,G_2)=\sum_{x\in\Sigma^*}\left(Pr_{G_1}(x)\log\frac{2Pr_{G_1}(x)}{Pr_{G_1}(x)+Pr_{G_2}(x)}+Pr_{G_2}(x)\log\frac{2Pr_{G_2}(x)}{Pr_{G_1}(x)+Pr_{G_2}(x)}\right)$$

The chi-squared (χ^2) *distance between two probabilistic grammars G_1 and G_2 is:*

$$d_{\chi^2}(G_1,G_2)=\sum_{x\in\Sigma^*}\frac{(Pr_{G_1}(x)-Pr_{G_2}(x))^2}{Pr_{G_2}(x)}$$

The Kullback-Leibler divergence, or relative entropy is not a metric:

Definition 10. *The KL divergence between two probabilistic grammars G_1 and G_2 is:*

$$d_{KL}(G_1,G_2) = \sum_{x\in\Sigma^*} Pr_{G_1}(x)\big(\log Pr_{G_1}(x) - \log Pr_{G_2}(x)\big) \tag{2}$$

Even if the KL-divergence does not respect the triangular inequality, $d_{KL}(G_1,G_2) = 0 \iff G_1 \equiv G_2$.

Definition 11. *Let G be a* PCFG. *The consensus string for G is the most probable string of $\mathscr{D}_G$. We denote by $\langle CS, \mathbf{PCFG(\Sigma)}\rangle$ the decision problem: is w the most probable string given G?*

2.5 PCP and the probabilistic grammars

Definition 12. *For each distance d_X the problem $\langle d_X, \mathscr{G}\rangle$ is the decision problem: given two grammars G and G' from $\mathscr{G}$, and any rational k, do we have $d_X(G,G') \leq k$?*

We use *Turing reduction* between decision problems and write:

$$\Pi_1 \leq_T \Pi_2$$

for problem Π_1 reduces to problem Π_2: if there exists a terminating algorithm solving Π_2 there also is one solving Π_1. If simultaneously $\Pi_1 \leq_T \Pi_2$ and $\Pi_1 \leq_T \Pi_2$, we will say that Π_1 and Π_2 are interreducible. The construction can be used for non decision problems: if only Π_1 is a decision problem, Π_1 is undecidable, and $\Pi_1 \leq_T \Pi_2$, we will say that Π_2 is *uncomputable*.

One particular well-known undecidable problem can be used as starting point for the reductions: the Post Correspondence Problem [31], which is undecidable:
Name: PCP
Instance: A finite set F of pairs of strings $(u_i, v_i), 1 \leq i \leq n$ over an alphabet Σ.
Question: Is there a finite sequence of integers $x_1 \dots x_t$, $t > 0$ such that $u_{x_1}u_{x_2}\dots u_{x_t} = v_{x_1}v_{x_2}\dots v_{x_t}$?

We give two standard constructions starting from an instance F of PCP. In both cases we use another alphabet containing one symbol $\#_i$ for each $i : 1 \leq i \leq n$. Let Ω denote this alphabet.

Construction 1: Two grammars
An instance of PCP as above is transformed into two PCFGs G_1 and G_2 as follows:
Rules of G_1: $S_1 \to u_i S_1 \#_i$ and $S_1 \to u_i \#_i$
Rules of G_2: $S_2 \to v_i S_2 \#_i$ and $S_2 \to v_i \#_i$
Each rule has probability $\frac{1}{2n}$.

Construction 2: One grammar
An instance of PCP is first transformed into two PCFGs G_1 and G_2 as above. Then a new non-terminal S_0 is introduced and we add the new rules $S_0 \to S_1$ and $S_0 \to S_2$, each with probability $\frac{1}{2}$.

The language obtained through G_0, G_1 and G_2 contains only strings x which can always be decomposed into $x = yz$ with $y \in \Sigma^*$ and $z \in \Omega^*$. We note that the number of derivation steps to generate string x is $1 + |z|$ for G_1 and G_2. For a final string x we denote this number by $\text{len}(x)$. For example $\text{len}(aababababa\#_3\#_1\#_4\#_1) = 4$.

Note that a positive instance of PCP will lead to G_1 and G_2 with a non empty intersection, and to an ambiguous G_0.

The following properties hold:

Property 1.

- G_1 *and* G_2 *are unambiguous and deterministic.*
- *If* $x \in \mathbb{L}(G_1)$, $Pr_{G_1}(x) = (\frac{1}{2n})^{\text{len}(x)}$
- F *is a positive instance of* PCP *if and only if* $\text{CoEM}(G_1, G_2) > 0$
- F *is a positive instance of* PCP *if and only if* G_0 *is ambiguous. In terms of probabilities, if there exists a string x such that* $Pr_{G_0}(x) = (\frac{1}{2n})^{\text{len}(x)}$

Property 2. *Let* F *be an instance of* PCP *with n pairs. Then*

$$\text{TAC}(G_0) = \frac{1}{16n-4}$$

$$\begin{aligned}
\text{TAC}(G_0) &= \sum_{i\geq 1}\left(n^i \cdot \frac{1}{4} \cdot (\frac{1}{2n})^{2i}\right) \\
&= \frac{n}{16n^2} \cdot \sum_{i\geq 0}(\frac{1}{4n})^i \\
&= \frac{1}{16n} \cdot \frac{1}{1-\frac{1}{4n}} \\
&= \frac{1}{16n} \cdot \frac{4n}{4n-1} \\
&= \frac{1}{16n-4}
\end{aligned}$$

3 Some decidability results

We report results concerning the problems related to equivalence and decision computation.

3.1 With the Manhattan distance

Proposition 3. *One cannot compute, given two* PCFG*s* G_1 *and* G_2*,*

$$d_{\text{L1}}(G_1, G_2) = \sum_{x\in\Sigma^*} |Pr_{G_1}(x) - Pr_{G_2}(x)|.$$

Proof. If this were possible, then we could easily solve the *empty intersection problem* by building $\text{MP}(G)$ and $\text{MP}(G')$ and then checking if $d_{\text{L1}}(\text{MP}(G), \text{MP}(G')) = 2$. □

Corollary 1. *One cannot compute, given two* PCFG*s* G_1 *and* G_2*, the variation distance between* G_1 *and* G_2*.*

The above follows from a straightforward application of Equation 1.

The same construction can be used for the Hellinger and Jenson-Shannon distances: whenever $\mathbb{L}(G_1)$ and $\mathbb{L}(G_2)$ have an empty intersection, it follows that $d_H(G_1, G_2) = 1$ and $d_{JS}(G_1, G_2) = 2$.

Summarizing, $\langle d_{\text{L1}}, \textbf{PCFG}(\Sigma)\rangle$, $\langle d_V, \textbf{PCFG}(\Sigma)\rangle$, $\langle d_H, \textbf{PCFG}(\Sigma)\rangle$ and $\langle d_{JS}, \textbf{PCFG}(\Sigma)\rangle$ are undecidable.

3.2 With the Euclidian distance

For PFA, positive results were obtained in this case: the distance can be computed, both for PFA and HMM in polynomial time[28].

In [19], Jagota *et al.* gave the essential elements allowing a proof that co-emission, the L2 and the Hellinger distances are uncomputable. In order to be complete, we reconstruct similar results in this section.

Proposition 4. *One cannot compute, given two* PCFG*s* G_1 *and* G_2*,*

$$\text{CoEm}(G_1, G_2) = \sum_{x \in \Sigma^*} Pr_{G_1}(x) \cdot Pr_{G_2}(x).$$

Proof. If this were possible, then we could easily solve the empty intersection problem, by taking G and G', building $\text{Mp}(G)$ and $\text{Mp}(G')$ and then checking if $\text{CoEm}(\text{Mp}(G), \text{Mp}(G')) = 0$. □

Proposition 5. *Computing the auto-coemission* AC *is at least as difficult as computing the* L2 *distance.*

Proof. Suppose we have an algorithm to compute the L2 distance. Then given any grammar G, we build a dummy grammar G_D which only generates, with probability 1, a single string over a different alphabet. It then follows that

$$d_{\text{L2}}(G, G_D) = \sqrt{\text{CoEm}(G,G) - 2\text{CoEm}(G,G_D) + \text{CoEm}(G_D, G_D)},$$

and since the intersection between the support languages for G and G_D is empty, $\text{CoEm}(G, G_D) = 0$. On the other hand $\text{CoEm}(G_D, G_D) = 1$, trivially. Therefore $\text{CoEm}(G,G) = d_{\text{L2}}(G, G_D)^2 - 1$. □

Corollary 2. *One cannot compute, given two* PCFG*s* G_1 *and* G_2*, the* L2 *distance between* G_1 *and* G_2*.*

Proof. By proposition 5 all we have to prove is that computing the auto-coemission is impossible. Let G_0 be the probabilistic context-free grammar built from an instance of PCP. Suppose we can compute $\text{AC}(G_0)$. Then since (by Property 2) we can compute $\text{TAC}(G_0)$, one could then solve the ambiguity problem via Proposition 2. This is impossible. □

Summarizing, $\langle d_{\text{L2}}, \textbf{PCFG}(\Sigma) \rangle$ and $\langle \text{CoEm}, \textbf{PCFG}(\Sigma) \rangle$ are undecidable. Furthermore $\langle \text{AC}, \textbf{PCFG}(\Sigma) \rangle$ and $\langle d_{\text{L2}}, \textbf{PCFG}(\Sigma) \rangle$ are interreducible.

3.3 With the KL divergence

Proposition 6. *One cannot compute, given two* PCFG*s* G_1 *and* G_2*,* $d_{KL}(G_1, G_2)$*.*

Proof. Suppose we could compute the KL divergence between two PCFGs. We should note that $d_{KL}(G_1, G_2) < \infty$ if and only if $\mathbb{L}(G_1) \subseteq \mathbb{L}(G_2)$. We would therefore be able to check if one context-free language is included in another one, which we know is impossible. □

The same proof can be used for the χ^2 distance since $d_{\chi^2}(G_1, G_2) < \infty$ if and only if $\mathbb{L}(G_1) \subseteq \mathbb{L}(G_2)$. Summarizing, $\langle d_{KL}, \textbf{PCFG}(\Sigma) \rangle$ and $\langle d_{\chi^2}, \textbf{PCFG}(\Sigma) \rangle$ are undecidable.

3.4 About the consensus string

A first result is that computing the Chebyshev distance is at least as difficult as computing the most probable string : Any PCFG G can be converted into G' with the same rules as G but using a disjoint alphabet. Now, it is clear that $d_{\mathrm{L}_\infty}(G,G') = Pr_G(CS(G))$:

Proposition 7. $\langle CS, \mathbf{PCFG(\Sigma)}\rangle$ *is decidable if there exists an algorithm computing* L_∞.
Ie, $\langle CS, \mathbf{PCFG(\Sigma)}\rangle \leq_T \langle d_{\mathrm{L}_\infty}, \mathbf{PCFG(\Sigma)}\rangle$.

Proposition 7 does not preclude that $\langle CS, \mathbf{PCFG(\Sigma)}\rangle$ may be decidable if $\langle \mathrm{L}_\infty, \mathbf{PCFG(\Sigma)}\rangle$ is not.
In fact $\langle CS, \mathbf{PCFG(\Sigma)}\rangle$ is decidable:

Lemma 1. *Let* $0 < \varepsilon < 1$. *Given any consistent* PCFG G, *there exists* $n \geq 0$ *such that* $Pr_G(\Sigma^{>n}) < \varepsilon$

Proof. Because $\lim_{n\to+\infty} Pr(\Sigma^{\leq n}) = 1$, there exists n such that $Pr(\Sigma^{\leq n}) \geq 1-\varepsilon$, hence $Pr(\Sigma^{>n}) < \varepsilon$. □

Proposition 8. $\langle CS, \mathbf{PCFG(\Sigma)}\rangle$ *is decidable.*

Proof. The proof for this is the existence of an algorithm that takes as input a PCFG G and always terminates by returning the consensus string for G. Algorithm 1 does exactly this.

□

Algorithm 1 goes through all the possible strings in Σ^0, Σ^1, Σ^2 ..., and checks the probability that G assigns to each string. It stores the string with the highest probability value (*Current_Best*) and the highest probability value itself (*Current_Prob*). It also subtracts from 1 all the probability values encountered (*Remaining_Prob*). So, after the i^{th} loop, *Current_Best* is the most probable string in $\Sigma^{<i}$, *Current_Prob* is the probability of *Current_Best* and *Remaining_Prob* is $1 - Pr(\Sigma^{<i})$ which is equal to $Pr(\Sigma^{\geq i})$. Using Lemma 1, we can say that for any ε, $0 < \varepsilon < 1$, there exists an i such that after the i^{th} iteration, *Remaining_Prob* is smaller than ε. This means that the algorithm must halt at some point. Moreover, if the most probable string in $\Sigma^{<i}$ has probability higher than $Pr(\Sigma^{\geq i})$, then we can be sure that this is the consensus string. This means that the algorithm always returns the consensus string.

3.5 With the Chebyshev distance

Proposition 9. $\langle d_{\mathrm{L}_\infty}, \mathbf{PCFG(\Sigma)}\rangle$ *and* $\langle EQ, \mathbf{PCFG(\Sigma)}\rangle$ *are interreducible.*

Proof. $\langle d_{\mathrm{L}_\infty}, \mathbf{PCFG(\Sigma)}\rangle \leq_T \langle EQ, \mathbf{PCFG(\Sigma)}\rangle$.
Suppose $\langle EQ, \mathbf{PCFG(\Sigma)}\rangle$ is decidable. Then if G_1 and G_2 are equivalent, $d_{\mathrm{L}_\infty}(G_1,G_2) = 0$. If G_1 and G_2 are not equivalent, there exists a smallest string x such that $Pr_{G_1}(x) \neq Pr_{G_2}(x)$. An enumeration algorithm will find this initial string x, whose probability is p. Note that if $Pr_{G_1}(\Sigma^{>n}) < p$ and $Pr_{G_2}(\Sigma^{>n}) < p$, we can be sure that no string in $\Sigma^{>n}$ has a difference of probabilities of more than p. This allows us to adapt Algorithm 1 to reach the length n at which we are sure that no string x longer than n can have probability more than $|Pr_{G_1}(x) - Pr_{G_2}(x)|$. The algorithm will therefore halt.

The converse ($\langle EQ, \mathscr{G}\rangle \leq_T \langle d_{\mathrm{L}_\infty}, \mathscr{G}\rangle$) is trivial since $d_{\mathrm{L}_\infty}(G_1,G_2) = 0 \Leftrightarrow G_1$ and G_2 are equivalent. □

```
Data: a PCFG G
Result: w, the most probable string
Current_Prob = 0;
Current_Best = λ;
Remaining_Prob = 1;
n = 0;
Continue = true;
while Continue do
    if Remaining_Prob < Current_Prob then
        Continue = false;
    else
        foreach w ∈ Σ^n do
            p = Pr_G(w);
            Remaining_Prob = Remaining_Prob − p;
            if p > Current_Prob then
                Current_Prob = p;
                Current_Best = w;
        n = n + 1;
return Current_Best
```

Algorithm 1: Finding the consensus string

4 Conclusion

The results presented in this work are threefold:

- the only positive result concerns the consensus string, which is computable;
- the multiple ambiguity problem (for CFGs) is equivalent to the Chebyshev distance problem (for PCFGs), which in turn is equivalent to the equivalence problem (also for PCFGs);
- all the other results correspond to undecidable problems.

Interestingly, if we consider the Chebyshev distance problem as a decision problem, namely:

Name: Chebyshev distance-$\leq$
Instance: Two PCFGs G_1 and G_2, $\varepsilon : 0 \leq \varepsilon \leq 1$
Question: Is $d_{\mathrm{L}_\infty}(G_1, G_2) \leq \varepsilon$?

the problem is actually decidable in all cases but one: when $\varepsilon = 0$. Ideally, one would hope to be able to bound the length of the strings over which the search should be done. This is possible in the case of probabilistic finite automata where it is proved that (1) a PFA can be transformed into an equivalent λ-free PFA $\mathscr{A}$, and, (2) the length of any string of probability at least p is upper bounded by $\frac{|\mathscr{A}|}{p^2}$, with $|\mathscr{A}|$ the size (number of states+1) of the PFA [17].

It should be noted that if the question is: Is $d_{\mathrm{L}_\infty}(G_1, G_2) < \varepsilon$?, the problem becomes decidable.

Moreover, it would be important to obtain approximation results, ie,

Name: X-distance-approx
Instance: Two PCFGs G_1 and G_2, $\varepsilon : 0 < \varepsilon \leq 1$

Question: Compute a such that $|a - d_X(G_1, G_2)| \leq \varepsilon$?

Such results have been studied in the case of probabilistic finite state machines, for example, recently, in [6]. In the case of the distances used in this work, the decidability of approximation would be ensured by Lemma 1. But the question of finding good approximations in polynomial time is clearly an interesting problem.

Acknowledgement

The authors thank James Worrell, Achilles Beros and Uli Fahrenberg for advice and discussions.

References

[1] S. Abney, D. McAllester & F. Pereira (1999): *Relating Probabilistic Grammars and Automata*. In: *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*.

[2] V. Balasubramanian (1993): *Equivalence and Reduction of Hidden Markov Models*. Master's thesis, Department of Electrical Engineering and Computer Science, MIT. Issued as AI Technical Report 1370.

[3] J-M. Benedí & J.-A. Sánchez (2005): *Estimation of stochastic context-free grammars and their use as language models*. *Computer Speech & Language* 19(3), pp. 249–274.

[4] R. C. Carrasco (1997): *Accurate computation of the relative entropy between stochastic regular grammars*. RAIRO *(Theoretical Informatics and Applications)* 31(5), pp. 437–444.

[5] F. Casacuberta & C. de la Higuera (2000): *Computational complexity of problems on probabilistic grammars and transducers*. In A. L. de Oliveira, editor: *Grammatical Inference: Algorithms and Applications, Proceedings of* ICGI *'00*, LNAI 1891, Springer-Verlag, pp. 15–24.

[6] T. Chen & S. Kiefer (2014): *On the total variation distance of labelled Markov chains*. In: *Proceedings of LICS 2014*, doi:`http://arxiv.org/abs/1405.2852`.

[7] C. Cortes, M. Mohri & A. Rastogi (2007): *L_p Distance and equivalence of probabilistic automata*. *International Journal of Foundations of Computer Science* 18(4), pp. 761–779.

[8] C. Cortes, M. Mohri, A. Rastogi & M. Riley (2008): *On the Computation of the Relative Entropy of Probabilistic Automata*. *International Journal on Foundations of Computer Science* 19(1), pp. 219–242.

[9] J. Esparza, A. Kucera & R. Mayr (2004): *Model Checking Probabilistic Pushdown Automata*. In: *Proceedings of LICS*, IEEE Computer Society, pp. 12–21.

[10] K. Etessami & M. Yannakakis (2009): *Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations*. *Journal of the ACM* 56(1), pp. 1–66.

[11] V. Forejt, P. Jancar, S. Kiefer & J. Worrell (2012): *Bisimilarity of Probabilistic Pushdown Automata*. *CoRR* abs/1210.2273.

[12] V. Forejt, P. Jancar, S. Kiefer & J. Worrell (2014): *Language Equivalence of Probabilistic Pushdown Automata*. *Information and Computation* 237, pp. 1–11.

[13] K. S. Fu & T. L. Booth (1975): *Grammatical Inference: Introduction and Survey. Part I and II*. IEEE *Transactions on Syst. Man. and Cybern.* 5, pp. 59–72 and 409–423.

[14] R. Gecse & A. Kovacs (2010): *Consistency of stochastic context-free grammars*. *Mathematical and Computer Modelling* 52(3–4), pp. 490 – 500.

[15] M. H. Harrison (1978): *Introduction to Formal Language Theory*. Addison-Wesley Publishing Company, Inc., Reading, MA.

[16] C. de la Higuera (2010): *Grammatical inference: learning automata and grammars*. Cambridge University Press.

[17] C. de la Higuera & J. Oncina (2013): *Computing the Most Probable String with a Probabilistic Finite State Machine*. In: *Proceedings of* FSMNLP.

[18] C. de la Higuera & J. Oncina (2013): *The Most Probable String: an Algorithmic Study*. *Journal of Logic and Computation, doi: 10.1093/logcom/exs049*.

[19] A. Jagota, R. B. Lyngsø & C. N. S. Pedersen (2001): *Comparing a Hidden Markov Model and a Stochastic Context-Free Grammar*. In: *Proceedings of* WABI *'01*, LNCS 2149, Springer-Verlag, pp. 69–84.

[20] F. Jelinek & J.D. Lafferty (1991): *Computation of the Probability of Initial Substring Generation by Stochastic Context-Free Grammars*. *Computational Linguistics* 17(3), pp. 315–323.

[21] F. Jelinek, J.D. Lafferty & R.L. Mercer (1992): *Basic Methods of Probabilistic Context Free Grammars*. In P. Laface & R. De Mori, editors: *Speech Recognition and Understanding — Recent Advances, Trends and Applications*, Springer-Verlag, pp. 345–360.

[22] M. Johnson (1998): *PCFG Models of Linguistic Tree Representations*. *Comput. Linguist.* 24(4), pp. 613–632.

[23] D. Jurafsky, C. Wooters, J. Segal, A. Stolcke, E. Fosler, G. Tajchaman & N. Morgan (1995): *Using a stochastic context-free grammar as a language model for speech recognition*. In: *Acoustics, Speech, and Signal Processing, 1995. ICASSP-95., 1995 International Conference on*, 1, IEEE, pp. 189–192.

[24] D. Klein & C. Manning (2003): *Accurate Unlexicalized Parsing*. In: *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics - Volume 1*, ACL '03, pp. 423–430.

[25] W. Kuich & A. Salomaa (1986): *Semirings, Automata, Languages*. Springer-Verlag.

[26] R. B. Lyngsø & C. N. S. Pedersen (2001): *Complexity of Comparing Hidden Markov Models*. In: *Proceedings of* ISAAC *'01*, LNCS 2223, Springer-Verlag, pp. 416–428.

[27] R. B. Lyngsø & C. N. S. Pedersen (2002): *The Consensus String Problem and the Complexity of Comparing Hidden Markov Models*. *Journal of Computing and System Science* 65(3), pp. 545–569.

[28] R. B. Lyngsø, C. N. S. Pedersen & H. Nielsen (1999): *Metrics and similarity measures for hidden Markov models*. In: *Proceedings of* ISMB *'99*, pp. 178–186.

[29] T. Murgue & C. de la Higuera (2004): *Distances between Distributions: Comparing Language Models*. In A. Fred, T. Caelli, R. Duin, A. Campilho & D. de Ridder, editors: *Structural, Syntactic and Statistical Pattern Recognition, Proceedings of* SSPR *and* SPR *2004*, LNCS 3138, Springer-Verlag, pp. 269–277.

[30] M-J. Nederhof & G. Satta (2004): *Kullback-Leibler distance between probabilistic context-free grammars and probabilistic finite automata*. In: *Proceedings of* COLING *'04 Proceedings of the 20th international conference on Computational Linguistics*, 71.

[31] E. L. Post (1946): *A variant of a recursively unsolvable problem*. *Bulletin of the American Mathematical Society* 52(4), pp. 264–268.

[32] Y. Sakakibara, M. Brown, R. Hughley, I. Mian, K. Sjolander, R. Underwood & D. Haussler (1994): *Stochastic context-free grammars for tRNA modeling*. *Nuclear Acids Research* 22, pp. 5112–5120.

[33] I. Salvador & J-M. Benedí (2002): RNA *Modeling by Combining Stochastic Context-Free Grammars and n-Gram Models*. *International Journal of Pattern Recognition and Artificial Intelligence* 16(3), pp. 309–316.

[34] A. Stolcke (1995): *An Efficient probablistic Context-Free Parsing Algorithm that Computes Prefix Probabilities*. *Computational Linguistics* 21(2), pp. 165–201.